arXiv:1608.06984v4 [cs.LG] 26 Apr 2017

Learning an Optimization Algorithm through Human Design Iterations

Thurston Sexton*[1] and Max Yi Ren†[1]

[1]Department of Mechanical Engineering, Arizona State University

April 28, 2017

Abstract

Solving optimal design problems through crowdsourcing faces a dilemma: On one hand, human beings have been shown to be more effective than algorithms at searching for good solutions of certain real-world problems with high-dimensional or discrete solution spaces; on the other hand, the cost of setting up crowdsourcing environments, the uncertainty in the crowd's domain-specific competence, and the lack of commitment of the crowd, all contribute to the lack of real-world application of design crowdsourcing. We are thus motivated to investigate a solution-searching mechanism where an optimization algorithm is tuned based on human demonstrations on solution searching, so that the search can be continued after human participants abandon the problem. To do so, we model the iterative search process as a Bayesian Optimization (BO) algorithm, and propose an inverse BO (IBO) algorithm to find the maximum likelihood estimators of the BO parameters based on human solutions. We show through a vehicle design and control problem that the search performance of BO can be improved by recovering its parameters based on an effective human search. Thus, IBO has the potential to improve the success rate of design crowdsourcing activities, by requiring only good search strategies instead of good solutions from the crowd.

*tbsexton@asu.edu

†yiren@asu.edu

1 Introduction

1.1 Challenges and opportunities for design crowdsourcing

Optimal design problems often have large solution spaces and highly non-convex objectives and constraints, inhibiting effective solution searching through existing optimization algorithms. Some of these problems, however, have been quite successfully (yet heuristically) solved by human beings. Notable examples include protein folding [1, 2], RNA synthesis [3, 4], genome sequence alignment [5], robot trajectory planning [6], and others [7–9]. The superior performance of *some* human beings at solving these problems demonstrates the advantages of human intelligence, which are supported by cognitive science and neuroscience findings [10] (see discussion in Sec. 5.1). However, despite a handful of success stories, applications of crowdsourcing to real-world design problems have yet to overcome several practical barriers. The cost of setting up problem-dependent crowdsourcing environments, the lack of commitment from crowd members, and uncertainty in domain-specific crowd competence have all contributed to its lack of adoption, while the growing availability of computation resources often makes straight-forward optimization or brute-force search a more convenient approach.

Our earlier study [8] highlighted these challenges for design crowdsourcing: We gamified a vehicle design and control problem (called the "ecoRacer" problem in what follows) where the objective is to complete a track with the minimal energy consumption within a time limit, by finding the optimal final drive ratio of the vehicle and the control policy for acceleration and regenerative braking. The game was broadcast on social media and received more than 2000 plays from 124 unique players within the first month. Results showed that (1) the marginal improvement in average game score of the crowd over an algorithm does not necessarily justify the high cost for developing crowdsourcing games, and (2) only a few players were committed to the search for more than 50 iterations, and still fewer can outperform the computer-found solution at all (see summary in Fig. 1).

Nonetheless, human search results displayed a significantly different *search pattern* than that of the algorithm. In particular, quite a few players showed rapid early improvement in performance, beyond the average performance of the computer, before they quit the game without reaching a solution close to the theoretical optimum. This observation is consistent with existing research (see, for example, [2] on a human-designed protein folding algorithm having a short-term advantage over a standard algorithm), and

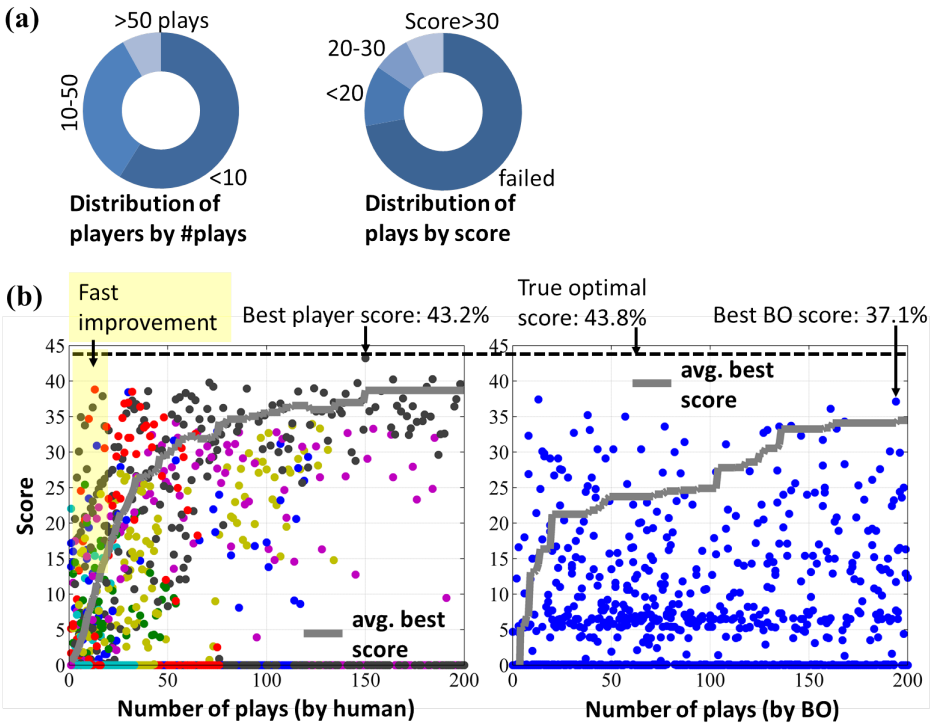

Figure 1: (a) Summary of player participation and performance (b) Results from the game showed while most players failed to outperform the Bayesian Optimization algorithm, some of them can identify good solutions early on. Image is reproduced from [8, 11].

suggests that while few people care to actually find the "best solution", their early demonstrations on *how they search* for a better solution may still be valuable. Specifically, we hypothesize that if a computer algorithm can be tuned to mimic these demonstrations, it can serve as a replacement to human solvers in their absence, to search in an effective way without ever abandoning the problem.

1.2 Learning to search

This paper aims to test the above hypothesis. We model a human solver's search behavior through a Bayesian Optimization algorithm (BO, also known as Efficient Global Optimization) [12, 13]. The algorithm iterates between two steps: (1) Estimating the shape of the problem space, based on previous solutions and corresponding performances, using a Gaussian Process (GP) model [14], and (2) creating a new solution based on this estimate (details in Sec. 2). While BO is not provably the underlying mechanism humans use, we hypothesize that the algorithm can be tuned to mimic the results of successful human search strategies, specifically in comparison with other popular gradient- and non-gradient-based optimization algorithms. The key assumption in modeling human search behavior through BO is the use of a GP to account for human beings' learning of input-output relationships (or called "function learning" in psychology). This assumption is supported by various findings: In a recent review of function-learning models, Christopher et al. [15] showed that the two major schools of models, i.e., rules- and similarity-based, can be unified through a Gaussian Process[1]. As discussed in Wilson et al. [16], the evidence that Occam's Razor plays an important role in human prediction also suggests that GP is an appropriate model for function learning, as GP reduces model complexity by construction [17]. Empirically, Borji et al. [18] showed that BO, with the use of GP, has the closest convergence performance to human searches when applied to 1D optimization problems. In fact, many higher-dimensional problems that human beings naturally solve, such as locomotion planning, have also been successfully solved through the use of GP [19–22].

Under this modeling assumption, we investigate how BO parameters can be estimated for the algorithm to best match human solver's search trajectory, i.e., the sequence of solution-performance pairs. To this end, we

[1]To be more accurate, the discussion in [15] is for function learning with continuous variables. While our case study involves discrete variables (acceleration and braking signals), the dimension reduction process converts these variables to continuous ones. See Sec. 4.

introduce an Inverse BO (IBO) algorithm to derive the maximum likelihood estimators for BO parameters, and discuss challenges in its implementation (see Sec. 3). Validation of the IBO algorithm takes two steps. We first use a simulation study to show that IBO can successfully estimate BO parameters used in generating a search trajectory (Sec. 3.2). We then show through the ecoRacer problem that the search performance of BO can be improved when its parameters are modified based on observing an effective human search and implementing IBO (Sec. 4). The results provide evidence that IBO can accelerate a search using only good search *strategies* without needing a large number of good human solutions. Thus, incorporating IBO in design crowdsourcing may lower the requirement on crowd commitment and so increase its chance of success. Limitations and their potential relaxations of the current IBO implementation will be discussed in depth in Sec. 5.

1.3 Related work

It is important to note that the focus of this paper is on the design of optimization algorithms aided by human demonstrations, rather than the derivation of qualitative explanations of the strengths and limitations of human design strategies. There have been numerous studies from the latter category in recent years (see [23–28] for example). This paper is also distinguished from studies that propose human-inspired optimization algorithms (see [29–31] for example), in that the learning of the optimization algorithm in our case is conducted by another algorithm, rather than by human researchers. From this aspect, our study is related to studies in learning-to-learn [32] where algorithms (e.g., for gradient-based optimization [33] and optimal control [34]) are tuned and controlled by a higher-level algorithm. In such work, however, the algorithms are often improved purely computationally through reinforcement learning by solving similar problems repeatedly. Due to the use of human demonstrations, our paper is also related to inverse reinforcement learning (see discussion in Subsec. 5.2), where human control strategies are used for defining and finding optimal control strategies.

2 Preliminaries on Bayesian Optimization

This section provides some background knowledge on BO to facilitate the discussion on IBO in Sec. 3.

2.1 Terminologies and notations

Let an optimization problem be $\min_{\mathbf{x}\in\mathcal{X}} f(\mathbf{x})$ where $\mathcal{X} \subseteq \mathbb{R}^p$ is the *solution space*. A search trajectory with K iterations can be represented by $h_K :=< \mathbf{X}_K, \mathbf{f}_K >$, where $\mathbf{X}_K$ and $\mathbf{f}_K$ represent the collection of K samples in $\mathcal{X}$ and their objective values, respectively. $h_0 :=< \mathbf{X}_0, \mathbf{f}_0 >$ represents an initial exploration set with K_0 samples. Human strategy is represented by algorithmic parameters $\boldsymbol{\lambda}$ that govern the search behavior: During the search, each new solution $\mathbf{x}_{k+1}$ (for $k = 0, \cdots, K-1$) is determined by $h_k :=< \mathbf{X}_k, \mathbf{f}_k >$ and $\boldsymbol{\lambda}$ through maximizing a merit function with respect to $\mathbf{x}$: $\mathbf{x}_{k+1} = \text{argmax}_{\mathbf{x}\in\mathcal{X}} Q(\mathbf{x}; h_k, \boldsymbol{\lambda})$. The functional forms of the merit function $Q(\mathbf{x})$ will be introduced in Subsecs. 2.2 and 3. We also define $\boldsymbol{\Lambda} := \text{diag}(\boldsymbol{\lambda})$ and its estimator as $\hat{\boldsymbol{\Lambda}} := \text{diag}(\hat{\boldsymbol{\lambda}})$.

2.2 The BO algorithm

We briefly review the BO algorithm, to explain how each new sample $\mathbf{x}$ is drawn based on the merit function $Q(\mathbf{x})$, itself defined by previous samples. Knowing this procedure is necessary for understanding the inverse BO algorithm, where we estimate the most likely BO parameters for a given trajectory of samples.

BO contains two major steps in each iteration: For a collection of samples of a black-box function, a Gaussian Process (GP) model is updated; the merit function is then formulated based on the GP model, and the next sample is chosen by maximizing the merit. **Model update**: It first updates a Gaussian Process (GP) model to predict objective values, based on current observations h_k and Gaussian parameters $\boldsymbol{\lambda}$. Without considering random noise in evaluating the objective, the GP model can be derived as $\hat{f}(\mathbf{x}; h_k, \boldsymbol{\lambda}) = b + \mathbf{r}^T\mathbf{R}^{-1}(\mathbf{f}_k - b)$, where $b = \frac{\mathbf{1}^T\mathbf{R}^{-1}\mathbf{f}_k}{\mathbf{1}^T\mathbf{R}^{-1}\mathbf{1}}$, $\mathbf{r}$ is a column vector with elements $r_i = \exp\left(-(\mathbf{x} - \mathbf{x}_i)^T\boldsymbol{\Lambda}(\mathbf{x} - \mathbf{x}_i)\right)$ for $i = 1, \cdots, k$, $\mathbf{R}$ is a symmetric matrix with $R_{ij} = \exp\left(-(\mathbf{x}_i - \mathbf{x}_j)^T\boldsymbol{\Lambda}(\mathbf{x}_i - \mathbf{x}_j)\right)$ for $i, j = 1, \cdots, k$, and $\mathbf{1}$ is a column vector with ones. Without prior knowledge, the Maximum Likelihood Estimator (MLE) of $\boldsymbol{\lambda}$ for the GP model can be derived by solving

$$\hat{\boldsymbol{\lambda}}_{GP} = \text{argmin}_{\boldsymbol{\lambda}} \log(\sigma^k |\mathbf{R}|^{\frac{1}{2}}), \tag{1}$$

where $\sigma^2 = (\mathbf{f}_k - \mathbf{1}b)^T\mathbf{R}^{-1}(\mathbf{f}_k - \mathbf{1}b)/n$ is the MLE of the GP variance. **Sampling the solution space**: The second step is to determine the next sample using the GP model. A common sampling strategy is to pick the new solution in $\mathcal{X}$ that maximizes the *expected improvement* from the current best objective value $f_{\min} := \min \mathbf{f}_k$ (assuming a minimization problem):

$Q_{EI}(\mathbf{x}; h_k, \boldsymbol{\lambda}) = (f_{\min} - \hat{f})\Phi\left(\frac{f_{\min}-\hat{f}}{\sigma}\right) + \sigma\phi\left(\frac{f_{\min}-\hat{f}}{\sigma}\right)$. Here $\Phi(\cdot)$ and $\phi(\cdot)$ are the cumulative distribution function and probability density function of the standard normal distribution, respectively. The new sample is thus obtained by solving

$$\mathbf{x}_{k+1} = \text{argmax}_{\mathbf{x}\in\mathcal{X}} Q_{EI}(\mathbf{x}; h_k, \boldsymbol{\lambda}). \tag{2}$$

Fig. 2 demonstrates four iterations of BO in optimizing a 1D function, with the GP model and the expected improvement function updated in each iteration. Note that similar to human searching behavior, BO is a stochastic process: First, the choice of the new design is stochastic, with better designs being more probable to be chosen[2]; and secondly, the initial exploration h_0 can be stochastic when it is modeled by a random sampling scheme, e.g., Latin Hypercube sampling (LHS, see [12] for details).

3 Inverse BO

We consider human solution search to consist two stages: A few exploratory searches are first conducted to acquire a preliminary understanding of the problem, before the execution of BO follows. For example, a player may spend a few trials to get familiar with a new game, before thinking about strategies to improve his score. IBO minimizes the sum of two costs corresponding to the exploration and BO stages, respectively. By doing so, it finds the most likely explanation of the underlying search strategy.

Specifically, IBO estimates $\boldsymbol{\lambda}$, along with the size of the initial exploration set K_0, given the trajectory h_K. To do so, we introduce and minimize a cost function consisting of the exploration cost for h_0, denoted as L_{INI}, and the BO cost for the rest of h_K, denoted as L_{BO}. We define $L_{INI} := -\log\left(Dp(\mathbf{X}_0)\right)$ where $p(\mathbf{X}_0)$ is the joint probability of the exploration set and $D := |\mathcal{X}|$ is the size of the solution space; and $L_{BO} := -\log\left(Dp(h_K - h_0|h_0)\right) = -\sum_{k=0}^{K-1}\log\left(Dp(\mathbf{x}_{k+1}|h_k)\right)$ where $p(\mathbf{x}_{k+1}|h_k)$ is the density for choosing $\mathbf{x}_{k+1}$ conditioned on h_k. Here $\log(\cdot)$ stands for natural logarithm.

The derivation of L_{INI} and L_{BO} are as follows: To calculate L_{INI}, we assume that each new sample during the exploration phase, $\mathbf{x}_i$ for $i = 1, \cdots, K_0$, tends to maximize its minimum Euclidean distance $d(\mathbf{x}_i, \mathbf{X}_{<i})$ to previous samples $\mathbf{X}_{<i}$, this is referred to as the *max-min sampling scheme*

[2]Numerically, this is because optimizing the non-convex function Q_{EI} requires a nested global optimization routine, such as Genetic Algorithm (GA), CMA-ES [35], DIRECT [36], and BARON [37]. Some implementations of these, e.g., GA and CMA-ES, can be stochastic.

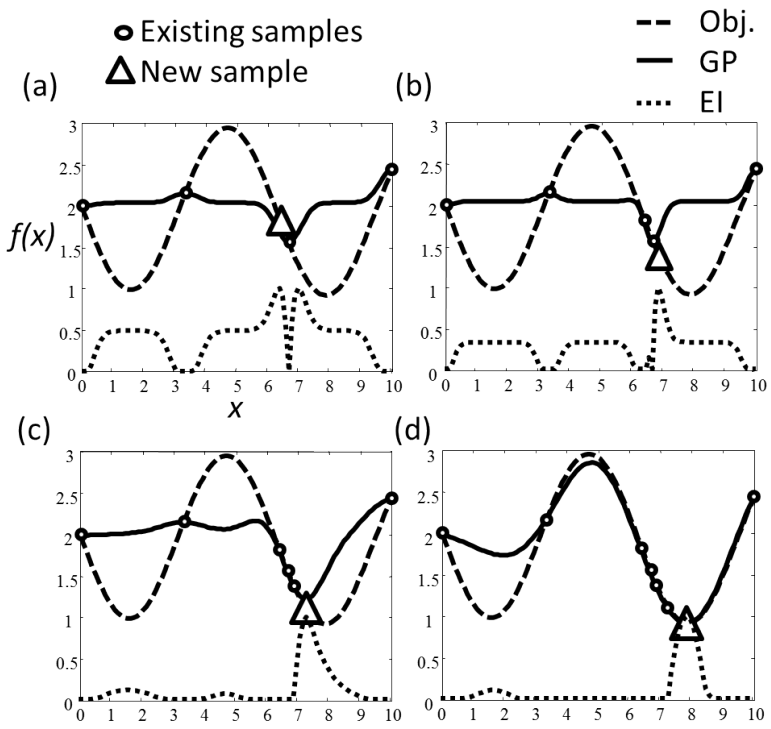

Figure 2: Four iterations of BO on a 1D function. Obj: The objective function. GP: Gaussian Process model. EI: Expected Improvement function. Image is modified from [11].

in what follows. Let the joint probability of the exploration set be $p(\mathbf{X}_0) = p(\mathbf{x}_1)p(\mathbf{x}_2|\mathbf{x}_1)\cdots p(\mathbf{x}_{K_0}|\mathbf{X}_{<K_0})$ and each conditional probability follow a Boltzmann distribution: $p(\mathbf{x}_i|\mathbf{X}_{<i}) = \exp\left(\alpha_{INI} d(\mathbf{x}_i, \mathbf{X}_{<i})\right) / Z_{INI}(\mathbf{x}_i, \alpha_{INI})$. Here the scalar α_{INI} represents how strictly each sample from $\mathbf{X}_0$ follows the max-min sampling scheme, and $Z_{INI}(\mathbf{x}_i, \alpha_{INI}) = \int_{\mathbf{x}\in\mathcal{X}} \exp\left(\alpha_{INI} d(\mathbf{x}, \mathbf{X}_{<i})\right) d\mathbf{x}$ is a partition function that ensures that $\int_{\mathcal{X}} p(\mathbf{x}_i|\mathbf{X}_{<i}) d\mathbf{x} = 1$. Note that the first sample in the exploration set is considered to be uniformly drawn, and thus its contribution to the cost (a constant) can be omitted.

To calculate L_{BO}, the conditional probability density of sampling $\mathbf{x} \in \mathcal{X}$ based on current h_k can be similarly modeled as a Boltzmann distribution:

$$p(\mathbf{x}|h_k) = \exp\left(\alpha_{BO} Q_{EI}(\mathbf{x}; h_k, \boldsymbol{\lambda})\right) / Z_{BO}(h_k, \boldsymbol{\lambda}, \alpha_{BO}), \tag{3}$$

where $Z_{BO}(h_k, \boldsymbol{\lambda}, \alpha_{BO}) = \int_{\mathbf{x}\in\mathcal{X}} \exp\left(\alpha_{BO} Q_{EI}(\mathbf{x}; h_k, \boldsymbol{\lambda})\right) d\mathbf{x}$ is also a partition function. The parameter α_{BO} plays a similar role to α_{INI}. For simplicity, we define $\tilde{l}_i := -\log\left(Dp(\mathbf{x}_i|\mathbf{X}_{<i})\right)$ and $l_k := -\log\left(Dp(\mathbf{x}_{k+1}|h_k)\right)$, so that $L_{INI} = \sum_{i=1}^{K_0} \tilde{l}_i$ and $L_{BO} = \sum_{k=0}^{K-1} l_k$. A lower value of $\tilde{l}$ or l represents higher probability density of the current sample to be drawn by max-min sampling or BO, respectively, and a zero indicates that the sample can be considered as uniformly drawn.

IBO solves the following problem to derive $\hat{\boldsymbol{\lambda}}$.

$$\min_{\alpha_{INI}, \alpha_{BO}, \boldsymbol{\lambda}, K_0} L := L_{INI} + L_{BO} \tag{4}$$

Note that to find the optimal K_0 for any given α_{INI}, α_{BO}, and $\boldsymbol{\lambda}$, one can first calculate the optimal $\tilde{l}_i$ and l_k for $i, k = 2, \cdots, K$, with respect to α_{INI}, α_{BO}, and $\boldsymbol{\lambda}$, and then scan $K_0 = 2, \cdots, K$ to find the lowest value of $L_{INI} + L_{BO}$. The scan starts at $K_0 = 2$ because it is not meaningful to initialize BO with a single sample.

3.1 Numerical Integration for Z_{BO}

The calculation of each l requires an approximation of the integral $Z_{BO}(h_k, \boldsymbol{\lambda}, \alpha_{BO})$, where the integrand $Q_{EI}(\mathbf{x}; h_k, \boldsymbol{\lambda})$ is usually a highly non-convex function with respect to $\mathbf{x}$, with function values dropping significantly around local maxima. See Fig.2 for example. Thus we propose to approximate Z_{BO} with importance sampling using a customized proposal density function that combines a uniform distribution with density $p(\mathbf{x}) = 1/D$ and a multivariate normal distribution with density $q(\mathbf{x}) = (\sqrt{2\pi}\sigma_I^p)^{-1} \exp(-||\mathbf{x} - \boldsymbol{\mu}||^2 / 2\sigma_I^2)$, where σ_I and $\boldsymbol{\mu}$ are parameters of $q(\mathbf{x})$. The uniform distribution is used to

sample over $\mathcal{X}$, while the normal distribution helps to improve the approximation by capturing the potential peak at the current sample $\mathbf{x}_{k+1}$. Thus we set $\boldsymbol{\mu} := \mathbf{x}_{k+1}$. Let $\mathbf{x}_i^u \in \mathcal{U}$ for $i = 1, ..., I$ and $\mathbf{x}_j^n \in \mathcal{N}$ for $j = 1, ..., J$ be samples from $p(\mathbf{x})$ and $q(\mathbf{x})$, respectively. The approximation $\hat{Z}_{BO}$ can be calculated by

$$\hat{Z}_{BO} := \sum_{\mathcal{U}} \frac{DQ_{EI}(\mathbf{x}_i^u)}{I\,(1 + Dq(\mathbf{x}_i^u))} + \sum_{\mathcal{N}} \frac{DQ_{EI}(\mathbf{x}_j^n)}{J\left(1 + Dq(\mathbf{x}_j^n)\right)}, \tag{5}$$

with arguments of Q_{EI} omitted for simplicity. The derivation of Eq. (5) is deferred to the appendix. Note that this approximation works under the assumption that $\int_{\mathbf{x}\in\mathcal{X}} q(\mathbf{x})d\mathbf{x} \approx 1$, which is plausible as the normal distribution is designed to have a narrow spread to match the local peak at $\mathbf{x}_{k+1}$. In this paper, the shape of this normal distribution is set by $\sigma_I = 0.01$ universally. While the setting of σ_I affects the variance of the approximation of Z_{BO}, we found this setting to perform well in practice. For Z_{INI}, since the minimum Euclidean distance function in a high dimensional space with limited samples is a relatively smooth function, we use Monte Carlo sampling for its approximation.

3.2 Simulation studies

As a validation step, we show that IBO can recover the parameters of a general BO given only an observed search trajectory. If IBO can determine the correct parameters (1) after a few number of iterations, (2) in a high-dimensional problem space, and (3) from a wide range of trajectory/parameter settings, then it could be used to recover parameters for matching a BO algorithm to an observed human search.

We use a simulation study to show that, for a given search trajectory, IBO can correctly identify the true $\boldsymbol{\lambda}$ provided the trajectory is sufficiently different from a random search. In addition, the simulation indicates that learning from already-efficient search behavior (i.e., estimating $\boldsymbol{\lambda}$ through IBO of an observed effective search trajectory) can lead to better BO convergence than the more common self-improvement methods (i.e., updating $\hat{\boldsymbol{\lambda}}$ by maximizing the likelihood of the observations according to the GP model).

3.2.1 Simulation settings and results

The simulation study is detailed as follows: We apply BO to a 30-dimensional Rosenbrock function constrained by $\mathcal{X} := [-2, 2]^{30}$. To initialize BO, we

use LHS to draw 10 samples from $\mathcal{X}$. BO terminates when the expected improvement for the next iteration is less than 10^{-3}. At each iteration, the expected improvement is maximized using a multi-start gradient descent algorithm [38] with 100 LHS initial guesses. A set of BO parameters, $\mathbf{\Lambda} = 0.01\mathbf{I}, 0.1\mathbf{I}, 1.0\mathbf{I}$, and $10.0\mathbf{I}$, are used to perform the search, where $\mathbf{I}$ is the identity matrix. For each of the four settings, 30 independent trials are recorded.

For each BO setting $\mathbf{\Lambda}$, each candidate estimator $\hat{\mathbf{\Lambda}}$, and each trajectory of length $K = 5, ..., 20$, we solve Eq. (4) using a grid search with $\mathcal{G}_{\alpha_{BO}} := \{0.01, 0.1, 1.0, 10.0\}$ and $\mathcal{G}_{K_0} := \{2, \cdots, K\}$. We fix α_{INI} to 1.0 and 10.0, and will discuss its influence to the estimation. Fig. 3 presents the resulting minimal L for all four cases and under all guesses. Each curve in each subplot shows how the minimal L (with respect to α_{BO} and K_0) changes as the search continues. The means and standard deviations of L are calculated using the 30 trials. Z_{INI} is approximated using a sample size of $10,000$. In approximating Z_{BO}, samples from the normal and the uniform distributions are of equal sizes ($I = J = 5,000$).

3.2.2 Analysis of the results

Based on the results from this simulation, as summarized in Fig. 3, the major finding from this simulation study is that IBO can successfully recover the BO parameters in cases where BO does not resemble uniform random sampling of the design space. In the cases of $\mathbf{\Lambda} = 0.01\mathbf{I}, 0.1\mathbf{I}, 1.0\mathbf{I}$, we see that the correct choices of $\hat{\mathbf{\Lambda}}$ consistently lead to the lowest cost along the search process. After only one or two iterations, in nearly all cases, the correct parameter has the highest likelihood of all four propositions, and this remains the case along the search. However, under large BO parameters such as $\mathbf{\Lambda} = 10.0\mathbf{I}$, the similarity between any two points in the design space becomes close to zero, leading to (almost) uniform uncertainty and expected improvement. Therefore this setting reduces BO to a uniform random sampling scheme. Fig. 3d shows that IBO does not perform well in this situation. To better understand the behavior of IBO under near-random searches, a curious reader may find a discussion on the properties of the costs l and $\tilde{l}$ in the Appendix.

3.2.3 Learning from others vs. self-adaptation

The above study showed that the correct BO setting $\boldsymbol{\lambda}$ can be learned through IBO. This subsection further demonstrates the advantage of "learn-

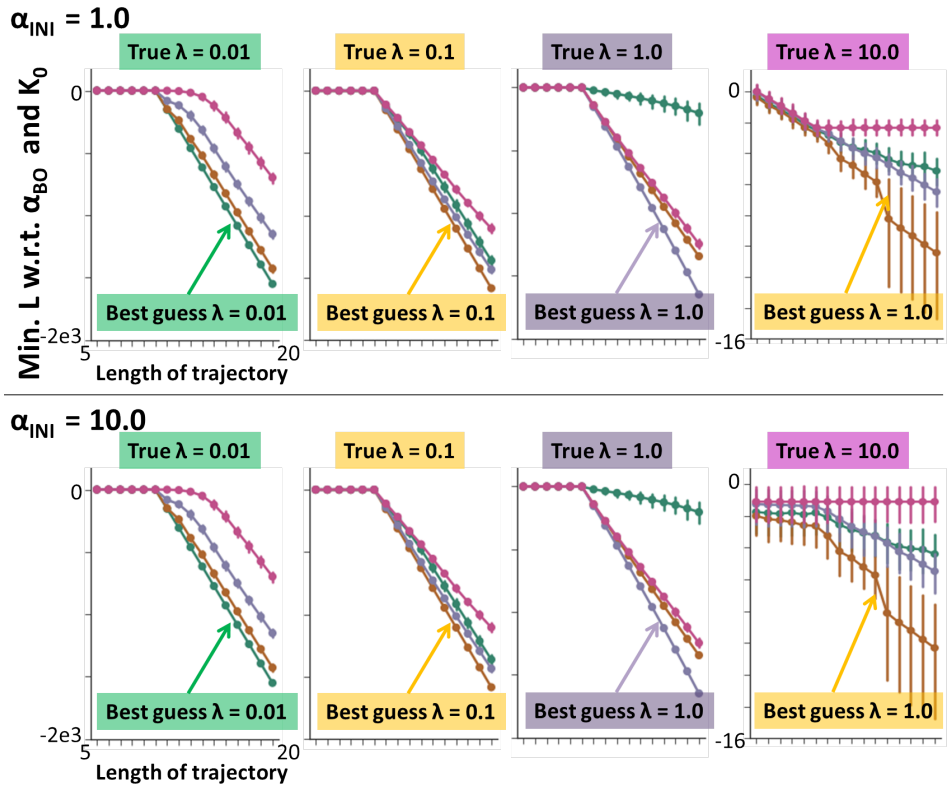

Figure 3: The minimal cost L for search trajectory lengths $N = 5, ..., 20$ with respect to $\mathcal{G}_{\alpha_{BO}}$ and $\mathcal{G}_{K_0}$. α_{INI} is fixed to 1.0 and 10.0. *View in color.*

ing from others" (i.e., updating $\boldsymbol{\lambda}$ through IBO), over "self-adaptation" (i.e., finding the MLE of $\boldsymbol{\lambda}$ using h_k). The settings follow the above study and results are shown in Fig. 4. First, to show the significant influence of $\boldsymbol{\lambda}$ on search effectiveness, we show the convergence of two fixed search strategies with $\boldsymbol{\Lambda} = 0.01$ and 10.0. Note that while neither converges to the optimal solution within 50 iteration, the former is significantly more effective than the latter. For "self-adaptive BO", we use a grid search ($\mathcal{G}_{\boldsymbol{\Lambda}} = \{0.01\mathbf{I}, 0.1\mathbf{I}, 1.0\mathbf{I}, 10.0\mathbf{I}\}$) to find $\hat{\boldsymbol{\Lambda}}_{\boldsymbol{GP}}$ that maximizes Eq. (1) at each iteration, and use $\hat{\boldsymbol{\Lambda}}_{\boldsymbol{GP}}$ to find the next sample. We show in Fig. 4b the percentages of the four guesses being $\hat{\boldsymbol{\Lambda}}_{\boldsymbol{GP}}$ along the search, using $\mathcal{G}_{\boldsymbol{\Lambda}}$ as the initial guesses for BO. The "learning from others" case starts with $\boldsymbol{\Lambda} = 10.0\mathbf{I}$ and uses IBO to derive $\hat{\boldsymbol{\Lambda}}$ from the trajectory produced by $\boldsymbol{\Lambda} = 0.01\mathbf{I}$. From Figs. 3 and 4b, we see that $\hat{\boldsymbol{\Lambda}}_{\boldsymbol{GP}}$ does not converge to $\boldsymbol{\Lambda} = 0.01\mathbf{I}$ as quickly as IBO, which explains why "learning from others" outperforms "self-adaptation" in Fig.4a. It is worth noting that this difference in performance may be relatively dependant on the dimensionality of the problem, as the two strategies were found to have similar convergence performance when applied to 2D functions. One potential explanation for this is that, in a lower dimensional space, an effective $\hat{\boldsymbol{\Lambda}}_{\boldsymbol{GP}}$ can be learned with a smaller number of samples.

4 Case study

We now investigate how IBO may improve the performance of BO when applied to a vehicle design and control problem.

4.1 Dimension reduction for player's control signals

The solution data from each game play consists of (1) the final gear ratio, (2) the recorded acceleration and braking signals, and (3) the corresponding game score. The length of a raw control signal matches that of the track, which has 18160 distance steps. Encoding control signals to a low dimensional space is feasible since common acceleration and braking patterns exist across all plays. In [11], this was done by introducing manually defined state-dependent basis functions (i.e., polynomials of the velocity of the car, slope of the track, distance to the terminal, remaining battery energy, and time spent) to parameterize the control signals. The underlying assumption that human players are aware of all the state-dependent bases is untested.

In this paper, we perform dimension reduction based on evidence that human beings often solve high dimensional problem by performing problem

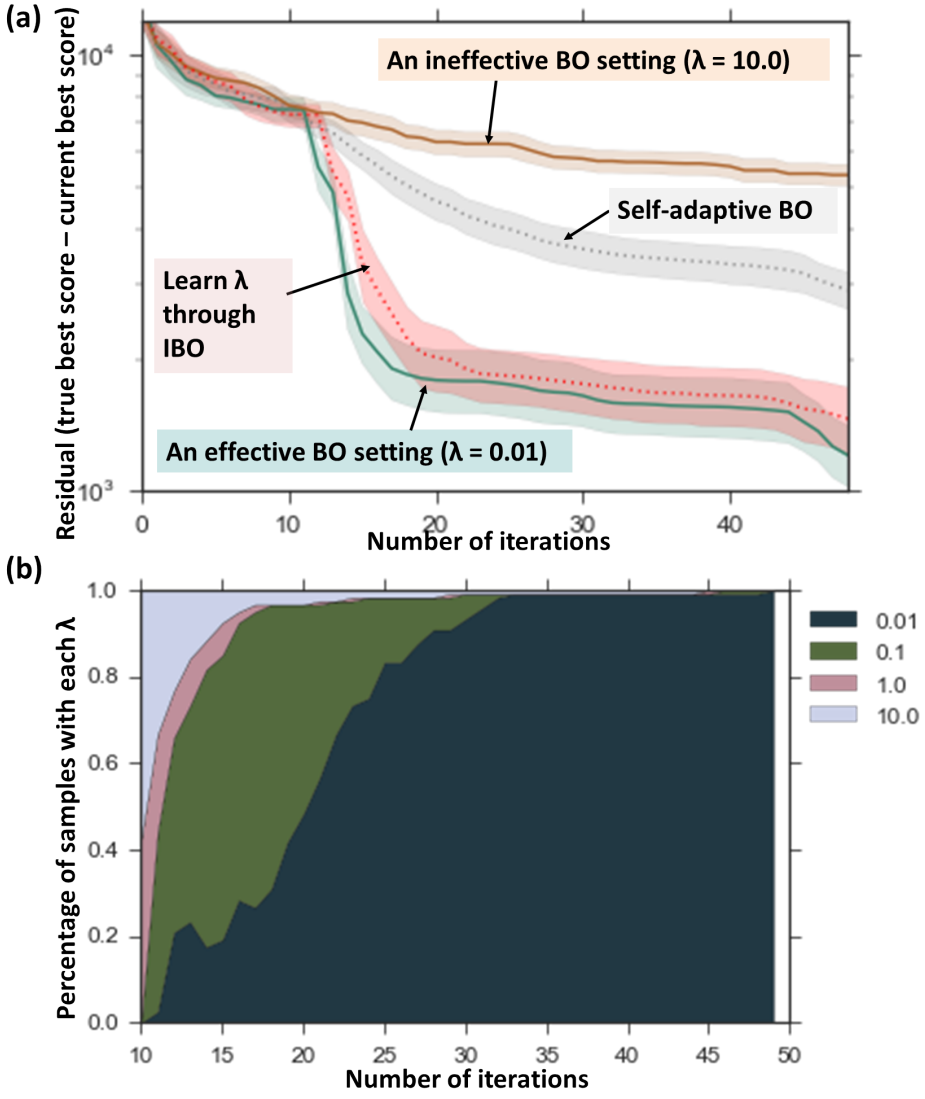

Figure 4: (a) Comparison on BO convergence using four algorithmic settings: (orange) $\mathbf{\Lambda} = 10.0\mathbf{I}$, (green) $\mathbf{\Lambda} = 0.01\mathbf{I}$, (grey) the MLE of $\mathbf{\Lambda}$ is used for each new sample, and (red) the initial setting $\mathbf{\Lambda} = 10.0\mathbf{I}$ is updated by IBO using the trajectory from $\mathbf{\Lambda} = 0.01\mathbf{I}$. (b) The percentages of estimated $\hat{\mathbf{\Lambda}}_{\boldsymbol{MLE}}$ along the number of iterations, averaged over the cases with $\mathbf{\Lambda} = \{0.01\mathbf{I}, 0.1\mathbf{I}, 1.0\mathbf{I}, 10.0\mathbf{I}\}$ and 30 trials for each case. View in color.

abstraction and using a hierarchical search [39–43]. In the context of the ecoRacer game, we hypothesize that players segment the track into m discrete sections, and make separate control decisions in each segment. Mathematically, this is equivalent to projecting observed signals onto m independent basis, which can be elegantly addressed by ICA [44]. Compared with Principal Component Analysis, where the bases minimize the covariance of the data, our ICA implementation maximizes the Kullback-Leibler divergence between all bases pairs, and is more suitable for non-Gaussian signals, such as the control data from this game (i.e., the acceleration/braking signals across players at each step along the track are unlikely to follow a Gaussian distribution).

Much like PCA, the choice of the number of ICA bases requires a balance between fidelity and practicality. While it is theoretically possible to find the "most likely" number of bases using information-theoretic criteria for model selection [45][3], we chose to use 30 bases because (1) over 95% of the variance is explained, and (2) the resultant solution space (30 control variables and one design variable) is small enough for BO to be effective.

4.2 Derivation of $\hat{\boldsymbol{\lambda}}$ and $\hat{\boldsymbol{\lambda}}_{GP}$

We apply IBO to two players, referred to as "P2" and "P3", who achieved the second and third highest score within 31 and 73 plays, respectively, much less than the 150 plays from the achiever of the highest score. To do so, we first encode all control solutions from the two players using the learned ICA bases. Together with the final drive ratios, all solutions are then normalized to be within $[-1, 1]^{31}$. IBO is performed separately on P2 and P3. We found that the probability for either player to have followed the max-min sampling scheme is lower than that of following BO, as the minimal values of $\tilde{l}(\mathbf{x}_k, \alpha_{INI})$ for $k = 2, ..., 31$ (with respect to α_{INI}) are dominated by those of $l(\mathbf{x}_k, \alpha_{BO})$. This means that the players were not likely to have performed an exploration before they started trying to improve their performance. This finding is reasonable, as the scoring mechanism in ecoRacer game, just like in other racer games with fairly predictable vehicle

[3]For completeness, we used 1000 PCA components as preprocessing to obtain the most likely number of ICA components under three suitable criteria: Minimum Description Length, Akaike Information Criterion, and Kullback Information Criterion, as 187, 464, and 373, respectively, using the method from [45]. While these dimensionalities could make sense from a neurological perspective (e.g., given that the game takes 36 seconds, a decision interval of 36s/187 = 192ms is close to the range for the time-frame of attentional blink, which is 200-500 ms [46]), the resultant high-dimensional solution spaces are unfavorable for BO.

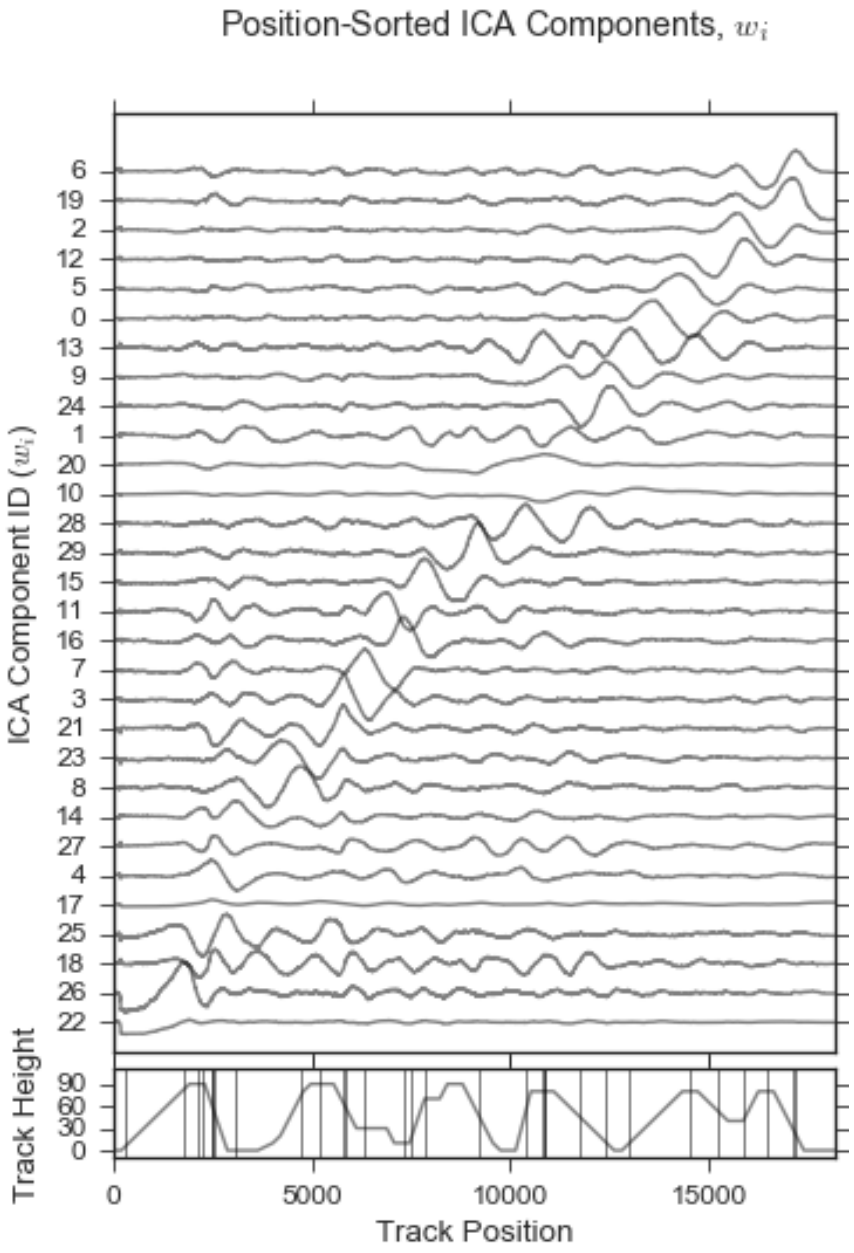

Figure 5: ICA bases learned from all human plays and the ecoRacer track. Vertical lines on the track correspond to the peak locations of the bases.

dynamics, can be understood by the player early on. Therefore, the search for $\hat{\boldsymbol{\lambda}}$ is performed by solving Eq. (4) with $\boldsymbol{\lambda} \in [0.01, 10.0]^{31}$, $\alpha_{BO} \in \mathcal{G}_{\alpha_{BO}}$, and a minimal number of initial samples ($K_0 = 2$) required for BO. For comparison purpose, we obtain $\hat{\boldsymbol{\lambda}}_{GP}$ using plays from P2, which represents a case where BO parameters are fine-tuned by the observed game plays, without trying to explain why these solutions were searched by the player.

Due to the non-convexity of Eq. (4) and Eq. (1), gradient-based searches using a series of 10 initial guesses are conducted to avoid inferior local solutions. Finite difference is used for gradient approximation. Both $\hat{\boldsymbol{\lambda}}$ and $\hat{\boldsymbol{\lambda}}_{GP}$ are calculated offline, and fixed during the execution of BO.

4.3 Comparison of BO performance

Fig. 6 compares the BO performance under $\hat{\boldsymbol{\lambda}}$ (for P2 and P3), $\hat{\boldsymbol{\lambda}}_{GP}$ and $\boldsymbol{\Lambda} = \mathbf{I}$. In each case, we start with the first two plays from the players, and run 180 BO iterations. Similar to the simulation study, results are reported using 20 trials due to the stochastic nature of BO. Due to the small trial number, bootstrap variance estimators are reported as the shades around the average in the figure. $\hat{\boldsymbol{\lambda}}$ outperforms the other two settings consistently along the search with statistical significance. The BO performance by mimicking P2 is slightly better than that of P3.

The result shows that BO can be improved noticeably by learning from P2 and P3. However, the players' search are not fully mimicked by IBO, as they improved much faster than the modified BO does, indicating that the proposed model still has room for improvement. Nevertheless, the IBO implementation still achieves the closest performance to the players' among all BO instances, and it is the only algorithm that achieved better performance than the players' best play within 100 iterations. This result demonstrates the potential of IBO to continue an effective human search after the player quits, with an improved search performance from a standard BO.

For completeness, we also note that in all cases, the BO identifies the true optimal final drive ratio at the end of the search. We also qualitatively compare the best human solution with one BO solution with high score, along with the theoretically optimal solution in Fig. 7. The result indicates that while these control strategies yield similar scores, they are quantitatively different, although braking towards the end is observed as a common strategy. Human search data are documented at `ecoracer.herokuapp.com/results`, where the best players' solution strategies are published.

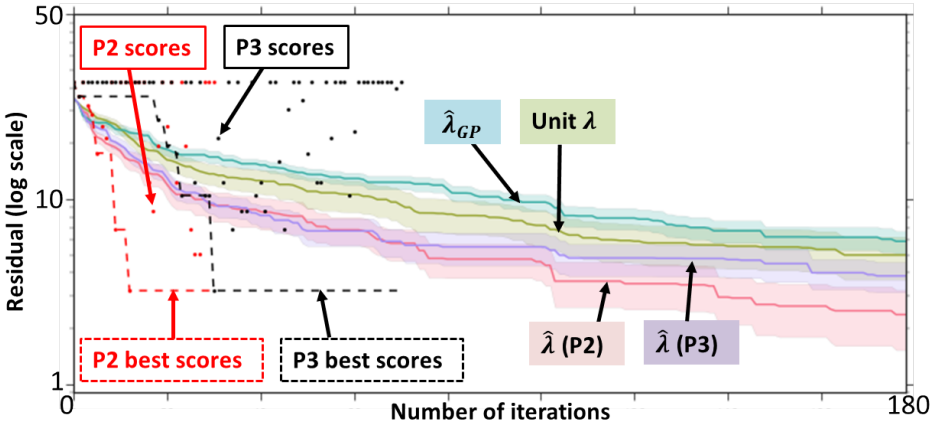

Figure 6: The residual of current best score vs. the known best score, with settings $\hat{\boldsymbol{\lambda}}$ (IBO, red), $\hat{\boldsymbol{\lambda}}_{GP}$ (MLE, blue), and the default $\boldsymbol{\lambda} = \mathbf{I}$ (green). Results are shown as averages over 30 trials. One-sigma confidence intervals are calculated via 5000 bootstrap samples. Red and black dots are scores from P2 and P3, respectively.

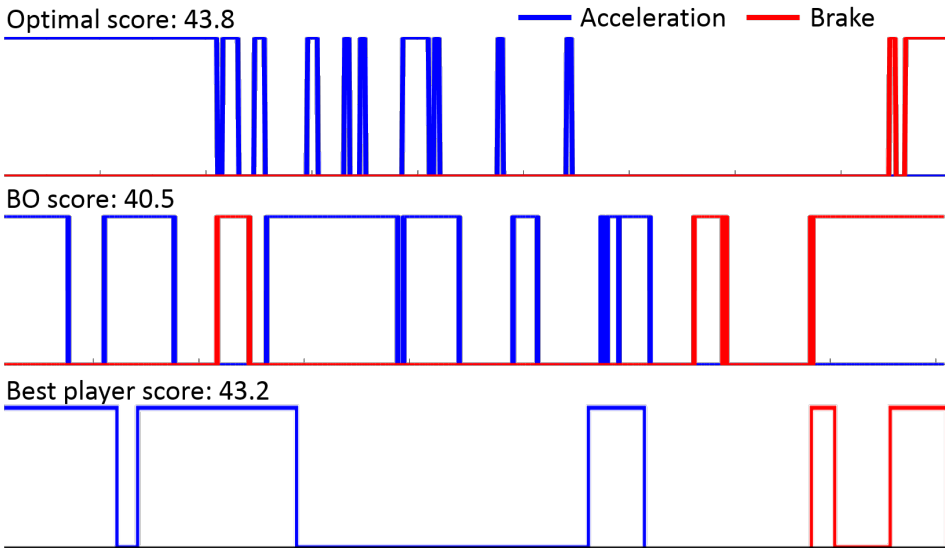

Figure 7: Qualitative comparison on control strategies from the theoretical optimal solution (top), one of the BO solutions (middle), and the best player solution (bottom).

5 Discussion

The above study provided a starting point for learning optimization algorithms based on human solution-search data. Yet, many pressing questions remain unanswered. This section will address a few notable ones. Some potential answers to these questions will rely on readers' familiarity with Inverse Reinforcement Learning [19, 47, 48] (IRL, also called apprenticeship learning [49, 50] and inverse optimal control [51]). To familiarize readers with this topic, a discussion on the connection between IBO and IRL is provided in Subsec. 5.2.

5.1 Limitations and potential values of IBO

From the case study, a strategy learning through IBO outperformed default algorithms, but is yet to reach the performance of the best human solver. This indicates potential room to further improve the algorithm. In the following, we discuss notable limitations of IBO. We shall also note that these also apply to the general problem of designing optimization algorithms through human demonstrations (called DO in what follows).

Model of human search strategies: Studies in cognitive science have put forth several core ingredients of human intelligence, including intuitive physics [52–55], problem decomposition skills [42, 56, 57], ability in learning-to-learn [58], and others [10]. While evidence has shown the connection between BO and human search [18], suitable models for human search strategies can be problem dependent. For example, for low-dimensional design problems, Egan et al. [59] showed that people adopting univariate search are more likely to achieve effective search. This result is supported by earlier psychological studies on how children perform scientific reasoning, and thus may be useful to explain how people identify unfamiliar systems. However, univariate search may not reflect how people search for solutions in a familiar context (such as car driving) and with a large number of control and design variables to tune, as is the situation of the ecoRacer game. For such high-dimensional and physics-based design and control problems, a potentially reasonable human search model could be to incorporate human intuitive physics models into the evaluation of the expected improvement. Thus instead of estimating GP parameters, one could estimate a statistical model of the state-space equations of the dynamical system, which influences the expected improvement. At a more abstract level, the fundamental challenge in understanding how a human search strategy should be modeled is the lack of knowledge about the functional form of the local objective

(i.e., the Q-function) that governs the generation of new solutions during the search based on the current state (cumulative knowledge learned by the human solver). As we will discuss later in this section, this challenge is also a key topic in IRL. Not surprisingly, one notable solution from IRL to this problem is in fact to use non-parametric models such as GP [19, 60].

Uncertainty in estimation: A limited amount of demonstrations could be insufficient to provide a good estimation of the BO parameters, even though the underlying parameters are the effective ones. One potential solution to this could be to create a reward mechanism in the crowdsourcing setting, where the reward is determined by both the observed search effectiveness of each human solver, and the uncertainty in the estimation of their search strategy. In the context of BO, this uncertainty can be measured by the covariance of the estimator, i.e., the Hessian of the cost function in Eq. (4). For people with effective search yet high estimation uncertainty, we can solicit more solutions from them by offering rewards. It would also be interesting to understand the influence of the properties of the problem, e.g., the size of the solution space, on the convergence of the estimation.

Knowledge transferability: The third limitation concerns the transferability of knowledge (search strategies) learned from one task (an optimization problem) to others. This limitation also leads to the question of how "effectiveness" of searches shall be measured, as we are not yet able to tell in what condition a strategy that has high rate of improvement (such as P2 in ecoRacer) will continue to produce better solutions than other strategies in a long term. The same issue, however, exists in IRL: e.g., a control policy learned for pancake flipping does not guarantee optimal egg flipping due to the differences in physical properties between pancakes and eggs. One solution to this in IRL is to allow the policy to adjust to new problem settings, by correcting the state transition model according to the new observations. This solution may also be applied to IBO. In the context of ecoRacer, knowledge such as "starting acceleration at the beginning of the track" could be considered as a universal strategy and requires less exploration, while the actual duration for executing this strategy may differ across problem settings. Therefore, it could be more effective for BO to adjust its parameters based on the ones that are learned from human demonstrations on a similar problem, rather than learning from scratch.

To summarize, IBO could be a valuable tool for machines to mimic human search behavior when (1) the underlying human search mechanism follows BO; (2) the demonstration is sufficient for estimating the true BO parameters with low variances, and (3) the true optimal BO parameters for a long-term search can be estimated based on an effective short-term search.

5.2 The difference between learning to search and learning a solution

The proposed IBO approach can be considered as a way to design optimization algorithms with human guidance, and is mathematically similar to IRL. In order to explain the similarities and differences between the two, we first introduce Markov Decision Process (MDP) and Reinforcement Learning (RL), and make an analogy between MDP and an optimization algorithm.

5.2.1 Preliminaries on MDP and RL

A MDP is defined by a tuple $< \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma, b_0 >$ where: $\mathcal{S}$ is a set of states; $\mathcal{A}$ is a set of actions; the state transition function $T(\mathbf{s}, \mathbf{a}, \mathbf{s}')$ determines the probability of changing from state $\mathbf{s}$ to $\mathbf{s}'$ when action $\mathbf{a}$ is taken; $R(\mathbf{s}, \mathbf{a})$ is the instantaneous reward of taking action $\mathbf{a}$ at state $\mathbf{s}$; $\gamma \in [0, 1)$ is the discount factor of future reward; $b_0(\mathbf{s})$ specifies the probability of starting the process at state $\mathbf{s}$. In RL, a control policy π is a mapping from a state to an action, i.e., $\pi : \mathcal{S} \to \mathcal{A}$. The long-term *value* of π for a starting state $\mathbf{s}$ can be calculated by $V^\pi(\mathbf{s}) = R(\mathbf{s}, \pi(\mathbf{s})) + \gamma \sum_{\mathbf{s}' \in \mathcal{S}} T(\mathbf{s}, \pi(\mathbf{s}), \mathbf{s}') V^\pi(\mathbf{s}')$, and thus the value of π over all possible starting states is the expectation $V^\pi = \sum_{\mathbf{s} \in \mathcal{S}} b_0(\mathbf{s}) V^\pi(\mathbf{s})$. A common way to represent a control policy is to introduce a Q-function $Q(\mathbf{s}, \mathbf{a}; \boldsymbol{\lambda})$ with unknown control parameters $\boldsymbol{\lambda}$, and let the policy be $\mathbf{a}(\mathbf{s}) = \text{argmax}_{\mathcal{A}} Q(\mathbf{s}, \mathbf{a}; \boldsymbol{\lambda})$. RL identifies the optimal $\boldsymbol{\lambda}$ that maximizes V^π.

5.2.2 MDP vs. optimization algorithm

An optimization algorithm defines a decision process: Its instantaneous reward is the improvement in the objective value achieved by each new sample, and the cumulative reward represents the total improvement in the objective within a finite number of iterations; its state contains the current solution (in $\mathcal{X}$), the corresponding objective value, and potentially the gradient and higher-order derivatives of the objective function at the current solution; its action is the next solution to evaluate; and its state transition is governed by the optimization algorithm and its parameters. This is similar to MDP where the state transition is affected by the control parameters. The decision process defined by an optimization algorithm, however, is usually non-Markovian, as the new solutions rely on the entire search trajectory. Note that it is still possible to consider the optimization process as an MDP, by redefining the state as the continuously growing search trajectory, i.e.,

elements in the state set $\mathcal{S}$ shall represent all possible search trajectories, rather than samples in $\mathcal{X}$.

5.2.3 IRL vs. IBO

RL algorithms identify an optimal control policy for an MDP with a given reward function. However, real-world applications hardly have explicit definitions of rewards, e.g., the reward for "driving a car" cannot be explicitly defined, although people form control policies based on their inherent reward (preference). Therefore, control policy for such applications can be learned more effectively through demonstrations of human beings, which are assumed to be optimal according to the inherent reward of the demonstrator. IRL techniques have thus been developed to identify the reward (and consequently the Q-function and the optimal control policy) that explains human demonstrations, either by estimating the reward parameters so that the demonstrated policy has a higher value than any other policies by a margin [47, 49, 61, 62], or by finding the maximum likelihood control parameters directly [48, 63].

The IBO approach introduced in this paper is closely related to latter type of IRLs, and more precisely, to the maximum entropy method of Ziebart et al. [48]. Briefly, the maximum entropy IRL proposes the following MLE of parameters $\boldsymbol{\lambda}$ based on a set of demonstrations h:

$$\begin{aligned}\hat{\boldsymbol{\lambda}} &= \operatorname{argmax}_{\boldsymbol{\lambda}} \log P(h|\boldsymbol{\lambda}) \\ &= \operatorname{argmax}_{\boldsymbol{\lambda}} \log \frac{\exp\left(\sum_{(\mathbf{s}_i,\mathbf{a}_i)\in h} R(\mathbf{s}_i,\mathbf{a}_i,\boldsymbol{\lambda})\right)}{\prod_{(\mathbf{s}_i,\mathbf{a}_i)\in h} Z_i(\boldsymbol{\lambda})},\end{aligned} \tag{6}$$

where $Z_i(\boldsymbol{\lambda})$ is a partition function for the visited state $\mathbf{s}_i$. One can notice the similarities between Eq. (6) and Eq. (4): (1) Both are maximum likelihood parameter estimations related to an instantaneous cost, i.e., the reward in Eq. (6) and the expected improvement in IBO. (2) Both involves partition functions that are computationally expensive, and dependent on the parameters $\boldsymbol{\lambda}$. Due to this dependency, a direct Markov-Chain Monte Carlo (MCMC) sampling in the space of $\boldsymbol{\lambda}$ (e.g., as in [63]) cannot be applied to optimize the likelihood function since the partition values for two different samples of $\boldsymbol{\lambda}$ do not cancel. Ziebart et al. discussed on alternative approach to address this computational challenge, by using the "Expected Edge Frequency Calculation" algorithm that has a complexity of $O(N|\mathcal{S}||\mathcal{A}|)$ for each gradient calculation of the objective in Eq. (6), where N is a large

number [48]. However, this approach can be infeasible for the IBO estimation problem in Eq. (4) since (1) the space $\mathcal{X}$ is usually continuous, and (2) even with a discretization of $\mathcal{X}$, the enormous size of $\mathcal{S}$ and $\mathcal{A}$ can easily make the calculation intractable, based on the discussion in Sec. 5.2.2.

Further, one shall notice that IRL and IBO uses different assumptions about human demonstrations: Demonstrations in IRL are assumed to be *near-optimal*. Thus learning from them leads to an optimal control policy for an MDP. Demonstrations in IBO, on the other hand, are assumed to be from an effect search strategy, yet are not necessarily optimal. Thus learning from them leads to an optimization algorithm, rather than a solution. This difference affects the application of the two: IRL can be used when the machine is told to mimic existing solutions, by understanding why these solutions are considered good, e.g., it answers the question "why do people flip pancakes this way?"; IBO can be used when the machine is meant to mimic the process of searching for good solutions, by understanding how to evaluate the expected improvement of solutions, e.g., it answers the question "how did people figure out this way of pancake flipping?".

6 Conclusions

In this paper, we attempted to address a dilemma in design crowdsourcing: While human beings acquire more advanced intelligence than machines in solving certain types of optimal design problems, soliciting valuable solutions through existing crowdsourcing mechanisms is not cost-effective due to the lack of control over crowd participation and the problem-specific qualification of the crowd. Based on the previous finding that more people acquire good searching strategies than good solutions, we proposed in this paper to mimic human search demonstrations by inversely learn a Bayesian optimization algorithm, so that long-term search can be executed more effectively by the computer even when human solvers abandon the problem. Through simulation and case studies, we showed improved performance of BO when it is equipped with parameters learned through an effective human search. However, the significant performance gap between a human demonstrator and the proposed algorithm in the case study suggested room for improvement of the algorithm. Future investigation will focus on closing this gap by exploring more suitable cognitive models of human solution searching for specific types of optimal design problems.

Acknowledgement

This work has been supported by the National Science Foundation under Grant No. CMMI-1266184. This support is gratefully acknowledged.

References

[1] Cooper, S., Khatib, F., Treuille, A., Barbero, J., Lee, J., Beenen, M., Leaver-Fay, A., Baker, D., Popović, Z., et al., 2010. Foldit. `http://fold.it`.

[2] Khatib, F., Cooper, S., Tyka, M. D., Xu, K., Makedon, I., Popović, Z., Baker, D., and Players, F., 2011. "Algorithm discovery by protein folding game players". *Proceedings of the National Academy of Sciences,* ***108***(47), pp. 18949–18953.

[3] Lee, J., Kladwang, W., Lee, M., Cantu, D., Azizyan, M., Kim, H., Limpaecher, A., Yoon, S., Treuille, A., and Das, R., 2014. eterna. `http://eterna.cmu.edu`.

[4] Lee, J., Kladwang, W., Lee, M., Cantu, D., Azizyan, M., Kim, H., Limpaecher, A., Yoon, S., Treuille, A., and Das, R., 2014. "Rna design rules from a massive open laboratory". *Proceedings of the National Academy of Sciences,* ***111***(6), pp. 2122–2127.

[5] Kawrykow, A., Roumanis, G., Kam, A., Kwak, D., Leung, C., Wu, C., Zarour, E., Sarmenta, L., Blanchette, M., Waldispühl, J., et al., 2012. "Phylo: a citizen science approach for improving multiple sequence alignment". *PloS one,* ***7***(3), p. e31362.

[6] Sung, J., Jin, S. H., and Saxena, A., 2015. "Robobarista: Object part based transfer of manipulation trajectories from crowd-sourcing in 3d pointclouds". *arXiv preprint arXiv:1504.03071*.

[7] Le Bras, R., Bernstein, R., Gomes, C. P., Selman, B., and Van Dover, R. B., 2013. "Crowdsourcing backdoor identification for combinatorial optimization". In Proceedings of the Twenty-Third international joint conference on Artificial Intelligence, AAAI Press, pp. 2840–2847.

[8] Ren, Y., Bayrak, A. E., and Papalambros, P. Y., 2016. "ecoracer: Game-based optimal electric vehicle design and driver control using human players". *Journal of Mechanical Design,* ***138***(6), p. 061407.

[9] Schrope, M., 2013. "Solving tough problems with games". *Proceedings of the National Academy of Sciences,* ***110***(18), pp. 7104–7106.

[10] Lake, B. M., Ullman, T. D., Tenenbaum, J. B., and Gershman, S. J., 2016. "Building machines that learn and think like people". *arXiv preprint arXiv:1604.00289*.

[11] Ren, Y., Bayrak, A. E., and Papalambros, P. Y., 2015. "ecoracer: Game-based optimal electric vehicle design and driver control using human players". In ASME 2015 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers, pp. V02AT03A009–V02AT03A009.

[12] Jones, D., Schonlau, M., and Welch, W., 1998. "Efficient global optimization of expensive black-box functions". *Journal of Global Optimization,* ***13***(4), pp. 455–492.

[13] Brochu, E., Cora, V. M., and De Freitas, N., 2010. "A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning". *arXiv preprint arXiv:1012.2599*.

[14] Rasmussen, C. E., 2006. "Gaussian processes for machine learning". MIT Press.

[15] Lucas, C. G., Griffiths, T. L., Williams, J. J., and Kalish, M. L., 2015. "A rational model of function learning". *Psychonomic bulletin & review,* ***22***(5), pp. 1193–1215.

[16] Wilson, A. G., Dann, C., Lucas, C., and Xing, E. P., 2015. "The human kernel". In Advances in Neural Information Processing Systems, pp. 2854–2862.

[17] Rasmussen, C. E., and Ghahramani, Z., 2001. "Occam's razor". *Advances in neural information processing systems*, pp. 294–300.

[18] Borji, A., and Itti, L., 2013. "Bayesian optimization explains human active search". In Advances in neural information processing systems, pp. 55–63.

[19] Levine, S., Popovic, Z., and Koltun, V., 2011. "Nonlinear inverse reinforcement learning with gaussian processes". In Advances in Neural Information Processing Systems, pp. 19–27.

[20] Deisenroth, M. P., Neumann, G., Peters, J., et al., 2013. "A survey on policy search for robotics.". *Foundations and Trends in Robotics,* ***2***(1-2), pp. 1–142.

[21] Calandra, R., Gopalan, N., Seyfarth, A., Peters, J., and Deisenroth, M. P., 2014. "Bayesian gait optimization for bipedal locomotion". In International Conference on Learning and Intelligent Optimization, Springer, pp. 274–290.

[22] Cully, A., Clune, J., Tarapore, D., and Mouret, J.-B., 2015. "Robots that can adapt like animals". *Nature,* ***521***(7553), pp. 503–507.

[23] Pretz, J. E., 2008. "Intuition versus analysis: Strategy and experience in complex everyday problem solving". *Memory & cognition,* ***36***(3), pp. 554–566.

[24] Linsey, J. S., Tseng, I., Fu, K., Cagan, J., Wood, K. L., and Schunn, C., 2010. "A study of design fixation, its mitigation and perception in engineering design faculty". *Journal of Mechanical Design,* ***132***(4), p. 041003.

[25] Daly, S. R., Yilmaz, S., Christian, J. L., Seifert, C. M., and Gonzalez, R., 2012. "Design heuristics in engineering concept generation". *Journal of Engineering Education,* ***101***(4), p. 601.

[26] Cagan, J., Dinar, M., Shah, J. J., Leifer, L., Linsey, J., Smith, S., and Vargas-Hernandez, N., 2013. "Empirical studies of design thinking: past, present, future". In ASME 2013 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers, pp. V005T06A020–V005T06A020.

[27] Björklund, T. A., 2013. "Initial mental representations of design problems: Differences between experts and novices". *Design Studies,* ***34***(2), pp. 135–160.

[28] Egan, P., and Cagan, J., 2016. "Human and computational approaches for design problem-solving". In *Experimental Design Research.* Springer, pp. 187–205.

[29] Cagan, J., and Kotovsky, K., 1997. "Simulated annealing and the generation of the objective function: a model of learning during problem solving". *Computational Intelligence,* ***13***(4), pp. 534–581.

[30] Landry, L. H., and Cagan, J., 2011. "Protocol-based multi-agent systems: examining the effect of diversity, dynamism, and cooperation in heuristic optimization approaches". *Journal of Mechanical Design,* ***133***(2), p. 021001.

[31] McComb, C., Cagan, J., and Kotovsky, K., 2016. "Drawing inspiration from human design teams for better search and optimization: The heterogeneous simulated annealing teams algorithm". *Journal of Mechanical Design,* ***138***(4), p. 044501.

[32] Thrun, S., and Pratt, L., 1998. "Learning to learn: Introduction and overview". In *Learning to learn*. Springer, pp. 3–17.

[33] Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., Blundell, C., Kumaran, D., and Botvinick, M., 2016. "Learning to reinforcement learn". *arXiv preprint arXiv:1611.05763*.

[34] Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M. W., Pfau, D., Schaul, T., and de Freitas, N., 2016. "Learning to learn by gradient descent by gradient descent". In Advances in Neural Information Processing Systems, pp. 3981–3989.

[35] Hansen, N., Müller, S. D., and Koumoutsakos, P., 2003. "Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es)". *Evolutionary computation,* ***11***(1), pp. 1–18.

[36] Jones, D. R., Perttunen, C. D., and Stuckman, B. E., 1993. "Lipschitzian optimization without the lipschitz constant". *Journal of Optimization Theory and Applications,* ***79***(1), pp. 157–181.

[37] Sahinidis, N. V., 1996. "Baron: A general purpose global optimization software package". *Journal of global optimization,* ***8***(2), pp. 201–205.

[38] Zhu, C., Byrd, R. H., Lu, P., and Nocedal, J., 1994. "L-bfgs-b: Fortran subroutines for large scale bound constrained optimization". *Report NAM-11, EECS Department, Northwestern University*.

[39] McGovern, A., Sutton, R. S., and Fagg, A. H., 1997. "Roles of macro-actions in accelerating reinforcement learning". In Grace Hopper celebration of women in computing, Vol. 1317.

[40] McGovern, A., and Barto, A. G., 2001. "Automatic discovery of subgoals in reinforcement learning using diverse density". *Computer Science Department Faculty Publication Series*, p. 8.

[41] Dietterich, T. G., 1998. "The maxq method for hierarchical reinforcement learning.". In ICML, Citeseer, pp. 118–126.

[42] Kulkarni, T. D., Narasimhan, K. R., Saeedi, A., and Tenenbaum, J. B., 2016. "Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation". *arXiv preprint arXiv:1604.06057*.

[43] Botvinick, M., and Weinstein, A., 2014. "Model-based hierarchical reinforcement learning and human action control". *Phil. Trans. R. Soc. B,* ***369***(1655), p. 20130480.

[44] Stone, J. V., 2004. *Independent component analysis*. Wiley Online Library.

[45] Hui, M., Li, J., Wen, X., Yao, L., and Long, Z., 2011. "An empirical comparison of information-theoretic criteria in estimating the number of independent components of fmri data". *PloS one,* ***6***(12), p. e29274.

[46] Tombu, M. N., Asplund, C. L., Dux, P. E., Godwin, D., Martin, J. W., and Marois, R., 2011. "A unified attentional bottleneck in the human brain". *Proceedings of the National Academy of Sciences,* ***108***(33), pp. 13426–13431.

[47] Ng, A. Y., Russell, S. J., et al., 2000. "Algorithms for inverse reinforcement learning.". In Icml, pp. 663–670.

[48] Ziebart, B. D., Maas, A. L., Bagnell, J. A., and Dey, A. K., 2008. "Maximum entropy inverse reinforcement learning.". In AAAI, pp. 1433–1438.

[49] Abbeel, P., and Ng, A. Y., 2004. "Apprenticeship learning via inverse reinforcement learning". In Proceedings of the twenty-first international conference on Machine learning, ACM, p. 1.

[50] Abbeel, P., Coates, A., and Ng, A. Y., 2010. "Autonomous helicopter aerobatics through apprenticeship learning". *The International Journal of Robotics Research*.

[51] Dvijotham, K., and Todorov, E., 2010. "Inverse optimal control with linearly-solvable mdps". In Proceedings of the 27th International Conference on Machine Learning (ICML-10), pp. 335–342.

[52] Spelke, E. S., Gutheil, G., and Van de Walle, G., 1995. "The development of object perception.".

[53] Baillargeon, R., Li, J., Ng, W., and Yuan, S., 2009. "An account of infants' physical reasoning". *Learning and the infant mind*, pp. 66–116.

[54] Bates, C. J., Yildirim, I., Tenenbaum, J. B., and Battaglia, P. W., 2015. "Humans predict liquid dynamics using probabilistic simulation". In Proceedings of the 37th annual conference of the cognitive science society.

[55] Gershman, S. J., Horvitz, E. J., and Tenenbaum, J. B., 2015. "Computational rationality: A converging paradigm for intelligence in brains, minds, and machines". *Science,* ***349***(6245), pp. 273–278.

[56] Fodor, J. A., 1975. *The language of thought*, Vol. 5. Harvard University Press.

[57] Biederman, I., 1987. "Recognition-by-components: a theory of human image understanding.". *Psychological review,* ***94***(2), p. 115.

[58] Harlow, H. F., 1949. "The formation of learning sets.". *Psychological review,* ***56***(1), p. 51.

[59] Egan, P., Cagan, J., Schunn, C., and LeDuc, P., 2015. "Synergistic human-agent methods for deriving effective search strategies: the case of nanoscale design". *Research in Engineering Design,* ***26***(2), pp. 145–169.

[60] Choi, J., and Kim, K.-E., 2012. "Nonparametric bayesian inverse reinforcement learning for multiple reward functions". In Advances in Neural Information Processing Systems, pp. 305–313.

[61] Ratliff, N. D., Bagnell, J. A., and Zinkevich, M. A., 2006. "Maximum margin planning". In Proceedings of the 23rd international conference on Machine learning, ACM, pp. 729–736.

[62] Syed, U., and Schapire, R. E., 2007. "A game-theoretic approach to apprenticeship learning". In Advances in neural information processing systems, pp. 1449–1456.

[63] Ramachandran, D., and Amir, E., 2007. "Bayesian inverse reinforcement learning". *Urbana,* ***51***, p. 61801.

Appendix

Derivation of Eq. (5)

Let $p(x) = 1/D$ and $q(x)$ be a uniform and a normal density function, respectively, D be the size of $\mathcal{X}$, and $f(x)$ be the function to be integrated. Also let $\mathcal{I}$ and $\mathcal{J}$ be the sample sets drawn from these two distributions, with sizes $I := |\mathcal{I}|$ and $J := |\mathcal{J}|$. We have

$$\begin{aligned}
\int f(x)dx &= D \int f(x)p(x)dx \\
&= D\left(\int \frac{f(x)p(x)^2}{p(x)+q(x)}dx + \int \frac{f(x)q(x)p(x)}{p(x)+q(x)}dx\right) \\
&\approx D\left(\frac{1}{I}\sum_{\mathcal{I}} \frac{f(x)p(x)}{p(x)+q(x)} + \frac{1}{J}\sum_{\mathcal{J}} \frac{f(x)p(x)}{p(x)+q(x)}\right) \\
&= \sum_{\mathcal{I}} \frac{f(x)D}{I(1+Dq(x))} + \sum_{\mathcal{J}} \frac{f(x)D}{J(1+Dq(x))}
\end{aligned}$$

IBO behavior under near-random search

Properties of l and $\tilde{l}$ From Sec. 3.1, the unbiased estimation of $l(\mathbf{x}, \alpha_{BO})$ through importance sampling is:

$$\hat{l}(\mathbf{x}, \alpha_{BO}) = -\log \frac{\exp(\alpha_{BO} Q_{EI}(\mathbf{x}))}{\hat{Z}_{BO}/D}. \tag{7}$$

$\hat{l}(\mathbf{x}, \alpha_{BO})$ has the following properties. **Property 1**: $\alpha_{BO} = 0$ leads to $\hat{l}(\mathbf{x}, 0) = 0$, indicating that $\mathbf{x}$ is uniformly sampled. One can see that the optimal cost of L_{BO} is non-positive, as one can always achieve $L_{BO} = 0$ by considering samples to be uniformly drawn. **Property 2**: When the expected improvement function is constant almost everywhere, i.e., $\Pr(Q_{EI}(\mathbf{x}) = C) = 1$, we have $\Pr(\hat{l}(\mathbf{x}, \alpha_{BO}) = 0) = 1$. This is because a uniformly drawn initial guess will almost surely satisfy the optimality condition for maximizing a constant function. **Property 3**: Notice that $1 + Dq(\mathbf{x}_i) \approx 1$ for $\mathbf{x}_i \in \mathcal{U}$ due to the small σ_I (see Sec. 3.1), and $\frac{\exp(\alpha_{BO} Q_{EI}(\mathbf{x}_i))}{1+Dq(\mathbf{x}_i)} \approx 0$ for large D and small α_{BO}. The partial derivative of $\hat{l}(\mathbf{x}, 0)$ with respect to α_{BO} can be approximated as:

$$\frac{\partial \hat{l}(\mathbf{x}, 0)}{\partial \alpha_{BO}} = c(\alpha_{BO}) \sum_{\mathcal{U}} \Delta a_i, \tag{8}$$

where $c(\alpha_{BO}) > 0$ and $\Delta a_i := Q_{EI}(\mathbf{x}_i) - Q_{EI}(\mathbf{x})$. Here we need to introduce a conjecture: Let $\bar{Q}_{EI} := \int_{\mathcal{X}} Q_{EI}(\mathbf{x}) d\mathbf{x}/D$ be the average expected improvement, and $A := \int_{\mathcal{X}} \mathbb{1}(Q_{EI}(\mathbf{x}) > \bar{Q}_{EI}) d\mathbf{x}$ be the measure of a subspace where the sampled expected improvement value is higher than $\bar{Q}_{EI}$. A decreases from above to below $D/2$ along the increase of the BO sample size. In other words, a uniformly drawn sample has more than 50% of chance to have an expected improvement value higher than $\bar{Q}_{EI}$ at the early stage of BO, and less than 50% at the late stage.

One evidence of the conjecture is illustrated in Fig. 2: In the first iteration, $\bar{Q}_{EI}$ is slightly lower than 0.5 while the majority of $\mathcal{X}$ has $Q_{EI} > \bar{Q}_{EI}$; in the fourth iteration, however, only a small region around the peak has $Q_{EI} > \bar{Q}_{EI}$. Using this conjecture, we can show that $\sum_{\mathcal{U}} \Delta a_i < 0$ when the sample size is small, thus $\frac{\partial \hat{l}(\mathbf{x},0)}{\partial \alpha_{BO}} < 0$. Together with Property 1, we have $\hat{l}(\mathbf{x}, \alpha_{BO}) < 0$ for a small α_{BO} and a small sample size.

Property 4: We notice that in this experiment, the discrepancy between LHS and the modeled max-min sampling scheme leads to overall high (positive) $\tilde{l}$ values, indicating that the samples are not likely follow this scheme. This is consistent with the fact that LHS is not exactly the same as max-min sampling, at least until all of h_0 has been considered. We also see that negative $\tilde{l}$ values can be observed when α_{INI} is low, suggesting that the LHS samples can be better explained by a loosely executed max-min sampling scheme than a strict one.

Discussion on findings from Fig. 3 We now summarize a complete list of findings based on these properties. **Finding 1**: A comparison between $\alpha_{INI} = 1.0$ and 10.0 leads to a finding consistent with Property 4. Since the samples are not likely to be drawn from a strictly executed max-min sampling scheme, the entire search trajectory is considered to be created from BO in the case of $\alpha_{INI} = 10.0$. While the early samples (less than 10) can be considered as from max-min sampling when $\alpha_{INI} = 1.0$ ($\tilde{l} < 0$), the low magnitude of $\tilde{l}$ causes this difference to be only visible in the case of $\mathbf{\Lambda} = 10.0\mathbf{I}$, where the magnitude of l is also low. **Finding 2**: IBO correctly identifies the true $\mathbf{\Lambda}$ within a few iterations after the initial exploration, except for the case of $\mathbf{\Lambda} = 10.0\mathbf{I}$. To explain this exception, we first note that $\mathbf{\Lambda} = 10.0\mathbf{I}$ leads to an expected improvement function that is constant almost everywhere (except for the sampled locations where $Q_{EI} = 0$) and thus BO reduces to uniform sampling. From Property 2, $L_{BO} = 0$ almost surely when we have the correct guess on $\mathbf{\Lambda}$. Also recall from Property 4 that $L_{INI} > 0$ when α_{INI} is high. The above two together explain why with

the correct guess of $\mathbf{\Lambda} = 10.0\mathbf{I}$, we have L close to zero when $\alpha_{INI} = 10.0$ and slightly negative when $\alpha_{INI} = 1.0$.[4]

To explain the negative L values for the incorrect guesses of $\mathbf{\Lambda}$, we use Property 3 to show that when the sample size is small and the expected improvement function is not flat, $L_{BO} < 0$ for a small α_{BO}, and thus $L < 0$. To summarize, Finding 2 suggests that for a search trajectory with a limited length that resembles a random search, the proposed IBO approach will consider it being derived from a BO that loosely solves Eq. 4. However, this caveat is of little practical concern, since (1) a random search rarely outperforms BO with non-trivial settings, and (2) a BO with low α_{BO} (instead of high $\mathbf{\Lambda}$) can equally simulate a random search.

[4]But why does the guess of $\mathbf{\Lambda} = 10.0\mathbf{I}$ lead to significantly decreasing L in the other three cases? This is because in those cases, BO does not resemble random sampling, i.e., the sequences of samples are more clustered. When a new sample is among this cluster, its similarities to existing ones are non-zero even when a large $\mathbf{\Lambda}$ is assumed, due to the small Euclidean distance among the pairs. And in turn, the expected improvement function has peaks within the clusters, and remains constant far away from them, rather than being a constant almost everywhere. As a result, the optimal value of $\hat{l}(\mathbf{x}, \alpha_{BO})$ with respect to α_{BO} becomes negative, even when $\mathbf{\Lambda}$ is incorrectly guessed as $10.0\mathbf{I}$.